EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	14	karl near dias.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:48
S2	7	shivani near gupta.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:49
S3	15	mark near ramacher.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:51
S4	6	uri near shaft.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:52
S5	4	venkateshwaran near venkataramani.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:53
S6	38	graham near3 wood.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:54
S7	58	juan near loaiza.in.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:57
S8	1496	oracle.as.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:57
S9	1	S8 and (data$1base and performance and solution).clm.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:58
S10	70	S8 and (data$1base and performance).clm.	US-PGPUB; USPAT	OR	ON	2007/03/27 16:58
S12	8	("2883255"|"20040044500"|"3324458"|"3351910"|"4654806"|"4994986"|"5488648"|"6954717").PN.	US-PGPUB; USPAT	OR	ON	2007/03/27 17:06
S13	3	("3344408"|"5349662"|"5872976").PN.	US-PGPUB; USPAT	OR	ON	2007/03/27 17:15
S14	1	("20040044700").PN.	US-PGPUB; USPAT	OR	ON	2007/03/27 17:18
S15	1	("6035306").PN.	US-PGPUB; USPAT	OR	ON	2007/03/27 17:18
S16	31	("4849879" | "5086386" | "5245638" | "5303166" | "5337258" | "5347647" | "5544313" | "5617567" | "5768500" | "5794229").PN. OR ("6035306").URPN.	US-PGPUB; USPAT; USOCR	OR	ON	2007/03/27 17:19
S17	3049	717/124-135.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:34

EAST Search History

S18	47	S17 and (data$1base and (performance adj problem))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:39
S19	47	S17 and (data$1base and (performance adj (problem or error)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:39
S20	1970	data$1base and (performance adj (problem or error))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:39
S21	680	S20 and ((diagnos$4 or analy$5 or examin$4) near3 performance)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:40
S22	78	S20 and (((diagnos$4 or analy$5 or examin$4) near3 performance) with data$1base)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:41
S23	52	S22 and (@pd<"20030905" or @ad<"20030905" or @prad<"20030905" or @rlad<"20030905")	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/27 17:41

EAST Search History

S24	99	("4994736" "5482046" "6021411" "6212528" "6212528" "6292830" "4780845" "5351197" "5459831" "5938736" "6072949" "6145028" "6247025" "5010478" "5265260" "5386583" "5412797" "5515513" "5530834" "5583986" "5631904" "5675797" "5676177" "5680618" "5721810" "5729466" "5734885" "5734903" "5737763" "5790805" "5799173" "5842208" "5860137" "5872795" "5887081" "5893074" "5895463" "5905868" "5913206" "5913207" "5917499" "5960423" "5978778" "5987453" "5991595" "6029231" "6058103" "6081801" "6169983" "6188992").pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/28 08:29
S25	5	("6615222" "6804627").pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/28 15:26
S26	2	"20020065833".pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/29 08:06

Terms used **database performance** Found **93,555** of **198,991**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Analysis of database performance with dynamic locking
In Kyung Ryu, Alexander Thomasian
July 1990 **Journal of the ACM (JACM)**, Volume 37 Issue 3
Publisher: ACM Press
Full text available: pdf(2.11 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A detailed model of a transaction processing system with dynamic locking is developed and analyzed. Transaction classes are distinguished on the basis of the number of data items accessed and the access mode (read-only/update). The performance of the system is affected by transaction blocking and restarts, due to lock conflicts that do not or do cause deadlocks, respectively. The probability of these events is determined by the characteristics of transactions and the database access pattern ...

2 Performance analysis of several back-end database architectures
Robert Brian Hagmann, Domenico Ferrari
March 1986 **ACM Transactions on Database Systems (TODS)**, Volume 11 Issue 1
Publisher: ACM Press
Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The growing acceptance of database systems makes their performance increasingly more important. One way to gain performance is to off-load some of the functions of the database system to aback-end computer. The problem is what functions should be off-loaded to maximize the benefits of distributed processing. Our approach to this problem consisted of constructing several variants of an existing relational database system. INGRES, that partition the database system software into tw ...

3 Locking performance in centralized databases
Y. C. Tay, Nathan Goodman, R. Suri
December 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 4
Publisher: ACM Press
Full text available: pdf(3.25 MB) Additional Information: full citation, abstract, references, citings, index terms

An analytic model is used to study the performance of dynamic locking. The analysis uses only the steady-state average values of the variables. The solution to the model is given by a cubic, which has exactly one valid root for the range of parametric values that is of interest. The model's predictions agree well with simulation results for transactions that require up to twenty locks. The model separates data contention from resource contention, thus facilitating an analysis of their separ ...

4 PYTHIA-II: a knowledge/database system for managing performance data and recommending scientific software
Elias N. Houstis, Ann C. Catlin, John R. Rice, Vassilios S. Verykios, Naren Ramakrishnan, Catherine E. Houstis
June 2000 **ACM Transactions on Mathematical Software (TOMS)**, Volume 26 Issue 2
Publisher: ACM Press
Full text available: pdf(796.18 KB) Additional Information: full citation, abstract, references, citings, index terms

Often scientists need to locate appropriate software for their problems and then select from among many alternatives. We have previously proposed an approach for dealing with this task by processing performance data of the targeted software. This approach has been tested using a customized implementation referred to as PYTHIA. This experience made us realize the complexity of the algorithmic discovery of knowledge from performance data and of the management of these data together with the d ...

Keywords: data mining, inductive logic programming, knowledge discovery in databases, knowledge-based systems, performance evaluation, recommender systems, scientific software

5 Transaction management for object-oriented databases: performance advantages of using multiple versions
Xiaoying Liu, John A. Miller, Nilesh R. Parate
April 1992 **Proceedings of the 25th annual symposium on Simulation ANSS '92**
Publisher: IEEE Computer Society Press
Full text available: pdf(985.78 KB) Additional Information: full citation, references, citings, index terms

6 Database performance in the real world: TPC-D and SAP R/3
Joachen Doppelhammer, Thomas Höppler, Alfons Kemper, Donald Kossmann
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2
Publisher: ACM Press
Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditionally, database systems have been evaluated in isolation on the basis of standardized benchmarks (e.g., Wisconsin, TPC-C, TPC-D). We argue that very often such a performance analysis does not reflect the actual use of the DBMSs in the "real world." End users typically don't access a stand-alone database system; rather they use a comprehensive application system, in which the database system constitutes an integrated component. In order to derive performance evalu ...

7 Florida International University High Performance Database Research Center
Naphtali Rishe, Wei Sun, David Barton, Yi Deng, Cyril Orji, Michael Alexopoulos, Leonardo Loureiro, Carlos Ordonez, Mario Sanchez, Artyom Shaposhnikov
September 1995 **ACM SIGMOD Record**, Volume 24 Issue 3
Publisher: ACM Press
Full text available: pdf(25.69 KB) Additional Information: full citation, citings, index terms

8 Empirical performance evaluation of concurrency and coherency control protocols for database sharing systems
Erhard Rahm
June 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 2
Publisher: ACM Press

Full text available: pdf(3.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a "shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar ...

Keywords: coherency control, concurrency control, database partitioning, database sharing, performance analysis, shared disk, shared nothing, trace-driven simulation

9 Research sessions: Research 14: Performance & tuning: Performance tradeoffs in read-optimized databases

Stavros Harizopoulos, Velen Liang, Daniel J. Abadi, Samuel Madden

September 2006 **Proceedings of the 32nd international conference on Very large data bases - Volume 32 VLDB'2006**

Publisher: VLDB Endowment

Full text available: pdf(735.37 KB) Additional Information: full citation, abstract, references, index terms

Database systems have traditionally optimized performance for write-intensive workloads. Recently, there has been renewed interest in architectures that optimize read performance by using column-oriented data representation and light-weight compression. This previous work has shown that under certain broad classes of workloads, column-based systems can outperform row-based systems. Previous work, however, has not characterized the precise conditions under which a particular query workload can be ...

10 The implementation and performance of compressed databases

Till Westmann, Donald Kossmann, Sven Helmer, Guido Moerkotte

September 2000 **ACM SIGMOD Record**, Volume 29 Issue 3

Publisher: ACM Press

Full text available: pdf(129.75 KB) Additional Information: full citation, abstract, citings, index terms

In this paper, we show how compression can be integrated into a relational database system. Specifically, we describe how the storage manager, the query execution engine, and the query optimizer of a database system can be extended to deal with compressed data. Our main result is that compression can significantly improve the response time of queries if very *light-weight* compression techniques are used. We will present such light-weight compression techniques and give the results of runni ...

11 Performance enhancements to a relational database system

Michael Stonebraker, John Woodfill, Jeff Ranstrom, Marguerite Murphy, Marc Meyer, Eric Allman

June 1983 **ACM Transactions on Database Systems (TODS)**, Volume 8 Issue 2

Publisher: ACM Press

Full text available: pdf(1.33 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we examine four performance enhancements to a database management system: dynamic compilation, microcoded routines, a special-purpose file system, and a special-purpose operating system. All were examined in the context of the INGRES database management system. Benchmark timings that are included suggest the attractiveness of dynamic compilation and a special-purpose file system. Microcode and a special-purpose operating system are analyzed and appear to be of more limited uti ...

Keywords: compiled query languages, database performance, file systems for databases, microcode

12 An analysis of database workload performance on simultaneous multithreaded processors

Jack L. Lo, Luiz André Barroso, Susan J. Eggers, Kourosh Gharachorloo, Henry M. Levy, Sujay S. Parekh

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.57 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

Simultaneous multithreading (SMT) is an architectural technique in which the processor issues multiple instructions from multiple threads each cycle. While SMT has been shown to be effective on scientific workloads, its performance on database systems is still an open question. In particular, database systems have poor cache performance, and the addition of multithreading has the potential to exacerbate cache conflicts.This paper examines database performance on SMT processors using traces of th ...

13 Integrating Database Technology with Comparison-based Parallel Performance Diagnosis: The PerfTrack Performance Experiment Management Tool

Karen L. Karavanic, John May, Kathryn Mohror, Brian Miller, Kevin Huck, Rashawn Knapp, Brian Pugh

November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

Publisher: IEEE Computer Society

Full text available: pdf(746.95 KB) Additional Information: full citation, abstract, index terms

PerfTrack is a data store and interface for managing performance data from large-scale parallel applications. Data collected in different locations and formats can be compared and viewed in a single performance analysis session. The underlying data store used in PerfTrack is implemented with a database management system (DBMS). PerfTrack includes interfaces to the data store and scripts for automatically collecting data describing each experiment, such as build and platform details. We have impl ...

14 Industry session 1: information retrieval: XML parsing: a threat to database performance

Matthias Nicola, Jasmi John

November 2003 **Proceedings of the twelfth international conference on Information and knowledge management CIKM '03**

Publisher: ACM Press

Full text available: pdf(210.86 KB) Additional Information: full citation, abstract, references, citings, index terms

XML parsing is generally known to have poor performance characteristics relative to transactional database processing. Yet, its potentially fatal impact on overall database performance is being underestimated. We report real-word database applications where XML parsing performance is a key obstacle to a successful XML deployment. There is a considerable share of XML database applications which are prone to fail at an early and simple road block: XML parsing. We analyze XML parsing performance an ...

Keywords: DOM, SAX, XML, database, parser, performance, validation

15 Data access performance in a large and dynamic pharmaceutical drug candidate database

Zina Ben-Miled, Yang Liu, Michael Bem, Robert Jones, Robert Oppelt, Samuel Milosevich, Dave Powers, Omran Bukhres

November 2000 **Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '00**

Publisher: IEEE Computer Society

Full text available: pdf(252.19 KB) Additional Information: full citation, abstract, references, index terms
Publisher Site

An explosion in the amount of data generated through chemical and biological experimentation has been observed in recent years. This rapid proliferation of vast amounts of data has led to a set of cheminformatics and bioinformatics applications that manipulate dynamic, heterogeneous and massive data. An example of such applications in the pharmaceutical industry is the computational process involved in the early discovery of lead drug candidates for a given target disease. This computationa ...

Keywords: cheminformatics, bioinformatics,multithreading, SMP, databases

16 Performance of a database manager in a virtual memory system

Stephen W. Sherman, Richard S. Brice
December 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 4
Publisher: ACM Press

Full text available: pdf(1.65 MB) Additional Information: full citation, abstract, references, citings, index terms

Buffer space is created and managed in database systems in order to reduce accesses to the I/O devices for database information. In systems using virtual memory any increase in the buffer space may be accompanied by an increase in paging. The effects of these factors on system performance are quantified where system performance is a function of page faults and database accesses to I/O devices. This phenomenon is examined through the analysis of empirical data gathered in a multifactor exper ...

Keywords: buffer manager, database management, double paging, page faults, page replacement algorithm, performance, virtual buffer, virtual memory

17 Implications of certain assumptions in database performance evauation

S. Christodoulakis
June 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 2
Publisher: ACM Press

Full text available: pdf(1.68 MB) Additional Information: full citation, abstract, references, citings, index terms

The assumptions of uniformity and independence of attribute values in a file, uniformity of queries, constant number of records per block, and random placement of qualifying records among the blocks of a file are frequently used in database performance evaluation studies. In this paper we show that these assumptions often result in predicting only an upper bound of the expected system cost. We then discuss the implications of nonrandom placement, nonuniformity, and dependencies of attribute ...

18 Parallelism and concurrency control performance in distributed database machines

Michael J. Carey, Miron Livny
June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data SIGMOD '89**, Volume 18 Issue 2
Publisher: ACM Press

Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references, citings, index terms

While several distributed (or 'shared nothing') database machines exist in the form of prototypes or commercial products, and a number of distributed concurrency control algorithms are available, the effect of parallelism on concurrency control performance has received little attention. This paper examines the interplay between parallelism and transaction performance in a distributed database machine context. Four alternative concurrency control algorithms are considered, including two-phas ...

19 A performance comparison of object and relational databases using the Sun

Benchmark
Joshua Duhl, Craig Damon
January 1988 **ACM SIGPLAN Notices , Conference proceedings on Object-oriented programming systems, languages and applications OOPSLA '88**, Volume 23 Issue 11
Publisher: ACM Press
Full text available: pdf(1.01 MB) Additional Information: full citation, abstract, references, citings, index terms

A general concern about object-oriented systems has been whether or not they are able to meet the performance demands required to be useful for the development of significant production software systems. Attempts to evaluate this assertion have been hampered by a lack of meaningful performance benchmarks that compare database operations across different kinds of databases. In this paper, we utilize the Sun Benchmark [Rube87] as a means for assessing the performance of an object d ...

20 Performance measurement methodologies for database systems
Steven A. Demurjian, David K. Hsiao, Douglas S. Kerr, Robert C. Tekampe, Robert J. Watson
October 1985 **Proceedings of the 1985 ACM annual conference on The range of computing : mid-80's perspective: mid-80's perspective ACM '85**
Publisher: ACM Press
Full text available: pdf(1.34 MB) Additional Information: full citation, references, index terms

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next